

PATENT APPLICATION
DOCKET NO. 72940/02-364

APPLICATION
Of
CHAD R. OVERTON, SUNIL MEHTA, JOHN M. LARSON
AND SCOTT E. ERRINGTON
For
UNITED STATES LETTERS PATENT
On
VIRTUAL MACHINE EMULATION IN THE MEMORY
SPACE OF A PROGRAMMABLE PROCESSOR

Docket: 72940/02-364
Sheets of Drawings: 10

Attorneys
Fellers, Snider, Blankenship, Bailey & Tippens, P.C.
Bank One Tower
100 North Broadway, Suite 1700
Oklahoma City, Oklahoma 73102-8820
Telephone: (405) 232-0621
Facsimile: (405) 232-9659

**VIRTUAL MACHINE EMULATION IN THE MEMORY SPACE
OF A PROGRAMMABLE PROCESSOR**

Related Applications

The present application claims domestic priority under 35 U.S.C. §119(e) to United States Provisional Patent Application No. 60/415,825 filed October 3, 2002.

5

Field of the Invention

This invention relates generally to the field of digital data processing devices and more particularly, but not by way of limitation, to a system and method for emulating a virtual machine within the memory space of a programmable processor using a virtual machine instruction set having a format that is independent of the hardware architecture of the processor, to allow the concurrent evaluation of execution of native code by the processor.

10

Background

15 Programmable processor-based devices are increasingly common in the current consumer products marketplace. One popular class of such devices are data storage devices, such as hard disc drives, optical disc (CD, DVD) players and recorders, magnetic digital tape (DAT) players and recorders, digital cameras, etc.

Such devices typically include a hardware/firmware based data interface circuit with a programmable processor (controller). A host device makes data transfer requests over a primary communication pathway using an established communications protocol (Fibre Channel, SCSI, ATA, RS-232, etc.).

20 The controller provides top level control of the operation of the device using associated programming stored in memory. Such programming includes operational routines that are executed as required to facilitate normal operation of the device, as well as test routines that can be selectively executed to certify and diagnose the operation of the device during manufacturing and field use. The test

25

routines can be accessed through the primary communication pathway or through a separate command pathway, such as a serial bus connected in parallel with the primary pathway.

Manufacturers often concurrently offer a number of different types and models of processor-based devices with varying levels of performance and features to meet different price points and other marketing requirements. Such devices may utilize different types and architectures of processors, such as digital signal processors, DSPs; standard PC-type processors such as the Intel Pentium® and X86 classes of processors; ARM devices (Advanced RISC (Reduced Instruction Set Computer) Machines), etc. Each of these generally has unique coding structures and requirements specific to the particular hardware architecture of the individual processors.

Because market pressures continue to press manufacturers to develop successive generations of processor-based devices with ever higher levels of performance at successively lower cost, there remains a continued need for improvements in the manner in which test code is generated and implemented, and it is to such improvements that the present disclosure is directed.

Summary of the Invention

In accordance with preferred embodiments, an apparatus and method are provided to emulate a virtual machine in the physical memory space of a programmable processor.

The apparatus generally includes a processor-based device in which native functions executed by the processor are stored in a format determined by a hardware architecture of the processor. The virtual machine is concurrently emulated by the processor to execute virtual functions each in a standardized format determined irrespective of the hardware architecture of the processor. Execution of a selected virtual function by the virtual machine causes the processor to execute a corresponding native function, with the corresponding native function being used to evaluate concurrent execution of at least one other native function.

Preferably, the virtual machine comprises a symbol table that maps each virtual function to the corresponding native function, an execution block coupled to the symbol table which executes the virtual functions, and a gate call interface

block which, when the execution engine block executes the selected virtual function, accesses the symbol table and initiates a gate call function to cause the processor to execute the corresponding native function.

Moreover, the execution of the corresponding native function by the
5 processor preferably results in a return data value, and the virtual machine further operates to provide the return data value to a host device. Preferably, the native functions comprise both operational routines and test routines, the operational routines carried out to facilitate normal operation of the device and the test routines carried out to identify error conditions associated with the normal operation of the
10 device. Hence, the virtual machine preferably operates to execute one of the native test routines to evaluate one or more of the native operational routines.

The method generally comprises steps of providing a set of virtual functions each in a standardized format determined irrespective of the hardware architecture of the processor, and using the virtual machine to execute a selected virtual function
15 which causes the processor to execute a corresponding native function, wherein the corresponding native function evaluates concurrent execution of at least one other native function.

The method further preferably comprises a prior step of a symbol table that maps selected virtual functions to the corresponding native functions, and the using
20 step further preferably comprises accessing the symbol table to identify the corresponding native function and performing a gate call operation to execute the corresponding native function.

In this way, virtual code in which the virtual functions are expressed can be established for a number of different processor architectures, facilitating
25 development, diagnostic and test functions by technical personnel in a fast and efficient manner.

These and various other features and advantages which characterize the claimed invention will be apparent from a reading of the following detailed description and a review of the associated drawings.

Brief Description of the Drawings

FIG. 1 is a top plan view of a processor-based data storage device constructed and operated in accordance with preferred embodiments of the present invention.

5 FIG. 2 is a functional block diagram of communication and control electronics of the device of FIG. 1.

FIG. 3 shows the device of FIG. 1 in conjunction with a host computer (PC).

10 FIG. 4 generally illustrates a computer network of a number of different data storage devices each utilizing a different processor architecture.

FIG. 5 provides a functional block diagram representing the memory space of a selected device controller (processor) which includes a virtual machine (VM) constructed and operated in accordance with preferred embodiments of the present invention.

15 FIG. 6 generally illustrates a preferred format for a VM register block shown in FIG. 5.

FIGS. 7 and 8 provide a list of preferred descriptions and formats of various VM functions of the VM code.

20 FIG. 9 is a functional block diagram illustrating a preferred methodology for generating the VM code.

FIG. 10 generally illustrates a preferred methodology for executing the VM code.

FIG. 11 generally illustrates the manner in which the symbol table is referenced during execution of the virtual machine.

25 FIG. 12 generally illustrates the manner in which the gate call interface (GCI) operates during execution of the virtual machine.

FIG. 13 is a flow chart for a VIRTUAL MACHINE EMULATION routine, generally representative of steps carried out in accordance with preferred embodiments of the present invention.

30

Detailed Description

To illustrate an environment in which presently preferred embodiments of the present invention can be advantageously practiced, FIG. 1 shows a processor-

based data storage device 100 of the type used to store and retrieve digital data. The device 100, preferably characterized as a hard disc drive, includes a base deck 102 which cooperates with a top cover 104 (shown in partial cutaway) to form an environmentally controlled housing for the device 100.

5 A spindle motor 106 supported within the housing rotates a number of rigid magnetic recording discs 108 in a rotational direction 109. An actuator 110 is provided adjacent the discs 108 and moves a corresponding number of heads 112 across the disc recording surfaces through application of current to an actuator coil 114 of a voice coil motor (VCM) 116. Data are written to and read from a number
10 of concentric tracks (not shown) defined on the disc surfaces by servo data written to the discs 108 during manufacturing. User data are stored on the tracks in data sectors (not shown) defined during a media formatting operation in areas between adjacent servo sectors.

FIG. 2 generally illustrates communication and control electronics of the
15 device 100 of FIG. 1. Interface (I/F) hardware 118 provides communication ports with the host device and includes a data buffer, error correction code (ECC) circuitry, and a sequencer (disc manager). Overall control is provided by a top level controller 120 with associated memory 122. The controller 120 and interface hardware 118 operate in accordance with a selected industry standard host
20 communications protocol, such as Fibre Channel.

A read/write application specific integrated circuit (R/W ASIC) 124 includes read and write channel circuitry to provide the requisite encoding for data to be written to the discs 108 during a write operation, and the requisite signal processing to decode transduced readback signals from the selected head 112
25 during a read operation. Such readback signals first undergo preamplification by a preamplifier/driver circuit (preamp) 126 mounted to the side of the actuator 110 (FIG. 1).

The R/W ASIC 124 includes servo demodulation capabilities that enable the device to decode the prerecorded servo data from the disc 108 and provide such
30 data to a digital ASIC 128. The digital ASIC 128 conditions the servo data by detecting the various different types of servo data (Gray code track address information, position dibit data, etc.) and arranging the data into multi-bit form (e.g., 16-bits).

The conditioned servo data are loaded into registers 130 for retrieval by a servo ARM (advanced RISC (reduced instruction set computer) machine) device 132. The ARM device 132 is a servo processor which operates in accordance with programming in ARM memory 134 and instructions from the controller 120. The
5 ARM device 132 outputs coil current commands to a coil driver 136 which in turn applies the appropriate current to the actuator coil 114 to position the selected head 112 as desired across the surface of the associated disc 108.

The digital ASIC 128 further includes a hardware manager 138 which provides various timing control functions for reference by the ARM device 132. It
10 will be noted that while the top level controller 120 and the servo ARM device 132 are shown as separate components, these separate functions can be carried out by a single processor device with sufficient bandwidth and processing capabilities to concurrently coordinate both data transfer and servo operations.

Normal data and control communications between the disc drive (HDD)
15 100 and a host device (host personal computer, PC) 140 are preferably carried out using a primary communications pathway (bus) 142, as shown in FIG. 3. In addition, a dedicated command communications pathway (serial bus) 144 is preferably utilized to allow access to test routines used to certify and/or diagnose the operation of the device 100 during manufacturing and field use. Alternatively,
20 some or all of the test routines can be accessed using the primary communications pathway 140.

FIG. 4 provides a generalized representation of a computer network 150 in a device manufacturing environment. A server PC 152 communicates with a number of host PCs 142, which are in turn locally connected to a corresponding
25 number of different devices (drives) 154, 156, 158, 160. The devices 154, 156, 158, 160 are generally similar to the device 100 discussed above, but utilize different top level and/or servo processors (denoted respectively as PROC A-D).

FIG. 4 serves to illustrate a challenge often faced by device manufacturers that provide a number of different types and models of devices that utilize different
30 processors with different hardware architectures. At a top level, the device manufacturer typically desires to carry out the same general types of certification and diagnostic tests on each type of device.

For example, it is contemplated that each type of device 154, 156, 158, 160 undergoes certain parametric configuration testing to set various parameters relating to the various circuits shown in FIG. 2 (R/W ASIC 124, digital ASIC 128, parameters and tables used by the servo ARM device 132, etc.).

5 The devices further typically undergo various types of certification type testing in which multiple test data patterns are written and reread to gauge error rate performance under various operational conditions (run-in), etc. These and other routines are performed to set up the devices for field use and to certify that the devices conform to various specifications prior to shipment. Such routines are
10 further advantageously employed during field use for diagnostic purposes.

 These and other routines are typically carried out by the top level controller 120 accessing specific test routines stored in top level controller memory 122. The routines are initiated by commands provided by the host PCs 140 via the primary or serial bus 142, 144, or executed independently on the drives internally. In a
15 high volume manufacturing environment such as contemplated by FIG. 4, it is desirable to automate such routines to reduce the need for human intervention. Thus, concurrent with the design effort to establish the design of each different device 154, 156, 158 and 160 is an effort to generate the requisite test routines to enable the certification and diagnosis of the devices.

20 Historically, a significant development effort has been necessary to generate the appropriate test routines in the native code of each processor, as well as the necessary initiation and control routines at the host PC 140 level to carry out the desired testing operations. As will be recognized by those skilled in the art, the development of code for a high performance, complex device is a structured
25 activity that involves numerous revisions and iterations as various improvements and modifications to other mechanical and electrical aspects of the device are implemented during the design process.

 At first glance, it might appear easier to simply wait and begin the test code development process until after the design is fully completed. However, such an
30 approach is impractical for several reasons. Not only would this approach add significant time to the overall design cycle, but more importantly, the test routines are often needed during the development process to help determine the viability and selection of various design alternatives along the way.

Another complicating factor is the existence and use of different revisions of code within the same type of product to support different interfaces, capacities, customer requirements, etc. This further increases the complexity of the diagnostic execution and reporting effort.

5 Accordingly, the present invention (as embodied herein and as claimed below) provides a significantly improved approach to the development and use of test (certification and diagnostic) routines. Generally, a virtual machine is developed as a standard processor with a standardized code set and library of operations. The virtual machine defines standard component interfaces which are
10 groups of functions that consistently represent particular hardware/software components of different devices regardless of the underlying architecture.

 The virtual machine is not realized in hardware as a discrete physical element, but instead is emulated within the memory spaces of each of a number of different processor platforms with different hardware architectures (e.g., PROC A-
15 D). Symbol tables and command gateways are utilized to translate standard commands into the native code instructions unique to each individual processor, and to format the output results into a standardized format for the user.

 FIG. 5 provides a schematic representation of a disc drive processor memory space 200. This memory space comprises the memory utilized by the
20 native processor of the device. For purposes of providing a concrete example, it will be contemplated that the memory space 200 corresponds to the memory 122 for controller 120 of the device 100 as shown in FIG. 2. However, it will be expressly understood that the same modules are generally provided to the respective memory spaces of a number of different types of devices, such as the
25 devices 154, 156, 158 and 160 of FIG. 4 (using the different processors PROC A-D).

 The memory space 200 includes a native code block 202 representative of the existing programming utilized by the controller 120. This native code comprises a sequence of native instructions in a format determined by the
30 hardware architecture of the controller 120. The controller 120 executes the native code to carry out various native functions of the device. The native functions thus represent subroutines of at least one instruction that carry out some desired operation. The native functions preferably include operational routines executed to

facilitate normal operation of the device 100 (such as the reading and writing of data) as well as test routines that operate to identify error conditions associated with the normal operation of the device 100 (such as the aforementioned certification and diagnostic functions).

5 The memory space 200 further includes a virtual machine (VM) block 204. Preferably, the VM block 204 is a 32-bit virtual machine that can be emulated by most modern processors. The VM block 204 provides complete support for structured programming while maintaining a small enough footprint to exist in an embedded environment. The VM block 204 is shown to preferably include a VM
10 execution engine 206, a gate call interface (GCI) block 208 and a symbol table 210. VM registers 212 and a VM server 214 are also provided. A VM code block 216 stores VM code to be executed by the engine 206, and a VM data block 218 stores return data values, such as variables, data structures, etc. The structure and operation of each will be discussed in detail below.

15 Generally, the engine 206 is a software state machine driven by the VM code and VM data via the VM registers 212. An exception to this is the use of gate call instructions (GCALL). When the engine 206 encounters a GCALL instruction, a native function is invoked via the GCI block 208. The GCI block 208 uses information in the symbol table 210 to map a virtual function with a
20 corresponding native function and to shepherd parameters and return values between the engine 206 and the native code 202.

 The VM block 204 is preferably designed to have the following characteristics:

1. Platform Independency: the VM architecture makes no assumptions
25 about the native processor Endian order, memory architecture, register architecture or instruction set; in other words, the VM code takes a format that is independent of the hardware architecture of the processor.

2. Native Optimization: the VM architecture does not prevent VM code from being converted into native instructions for faster execution.

30 3. Structured Programming: the VM architecture provides the necessary instruction set to support structured programming. Preferably, C language is the base line.

4. Native Function Access: native functions that are accessed by VM code are not exclusive for the VM, but are rather available to the native code without special consideration.

5. Safety: the VM architecture isolates VM code and data from native code and data. The VM code can preferably only invoke those native functions that are made available to it through the GCI block.

For reference, it will be helpful to define the term “VM code” to preferably describe a block of 16-bit values that tell the VM block 204 what operation to perform. The term “VM data” is preferably defined as a block of 8-bit values that are referenced as any combination of values or arrays of values associated with the particular operation in progress.

The engine 206 is invoked with three non-overlapping native data regions containing the VM code, VM data and VM registers (i.e., blocks 216, 218 and 214, respectively). Each time that the engine 206 access the VM code or VM data, it verifies that those accesses are contained within the regions originally specified at invocation. Consequently, the engine 206 operates within “sandboxes” to prevent errant or malicious code from adversely affecting operation.

The VM instruction set is a modified load/store architecture. Three classes of data are recognized: global, local and parameter. This architecture limits the number of instructions that require checking sandboxes, and the designation of these three classes of data reduce the need for expensive stack manipulations required to access stack-based local variables and parameters. A preferred listing and description of various instructions is provided hereinbelow.

As shown in FIG. 6, the VM registers 212 preferably comprise eight directly accessible 32-bit registers (denoted R0-R5, SP and IP) and an indirectly accessible 32-bit CPU Status Register (CPUSR). The contents of the directly accessible registers can be explicitly manipulated via the instruction set, whereas the contents of the CPUSR register are set implicitly based on the outcome of certain VM instructions. The eight directly accessible registers can be accessed as the lower byte (8-bit), lower word (16-bit) or the entire double-word (32-bit). The CPUSR register is preferably used primarily for conditional branches and complex mathematical operations.

The general purpose registers R0-R5 are commonly used to provide the returned value from a called function, or register based local variables. Otherwise, there are generally no restrictions on the use of these registers.

5 The SP (stack pointer) register points to the lowest address element on the stack and is used by the PUSH, POP and CALL instructions. The VM Stack operates as a decrement-and-store for a PUSH instruction and a load-and-increment for a POP instruction. An implicit PUSH IP occurs during a register based CALL instruction.

10 The IP (instruction pointer) register points to the next instruction to be processed. The lowest bit of this register is preferably ignored. Each instruction cycle begins with a fetch of the instruction pointed to by the IP register and ends with the incrementing of the IP register by 2. The combination of the CALL instruction and the POP IP instruction create the call/return sequence used in sub-routine handling.

15 The CPUSR register preferably contains four operational flags: zero, carry, negative and overflow (ZCNV), and three control flags: abort, branch and debug (ABD). The flags occupy bits 31-28 and bits 2-0, respectively. To optimize performance, not all instructions utilize the flags as such are not always necessary for every type of instruction.

20 The VM code and data are addressed separately and are generally not allowed to overlap. VM code is read-only and is implicitly accessed according to the value in the IP register at the beginning of each instruction cycle. VM data are accessed only through the load and store instructions (LDG, STG, LDL, STL, PUSH and POP, discussed below).

25 VM code elements are 16-bit quantities and are aligned on two-byte boundaries. All VM instructions are preferably encoded as one or more of these 16-bit quantities. VM data elements preferably have no alignment restrictions and can be addressed as quantities of one, two or four bytes.

30 VM code and data are accessed by the engine 206 in the native Endian Order of the particular implementation associated with the particular processor (i.e., whether lower byte addresses at a given location hold the least significant bits or the most significant bits). Any process that places VM code and data into memory is responsible for ensuring that the correct Endian Order is utilized prior

to invoking the engine 206. Generally, programmers can determine Endian Order scheme at compiling time, if such information is not already known for the given platform. For example, a 2-byte value such as 0x55AAH can be written to a selected location. When a 1-byte value is subsequently read back from the same location, a readback value of 0x55H will indicate a Big-Endian platform while a readback value of 0xAAH will indicate a Little-Endian platform.

The GCI block 208 operates as a link between VM functions and native functions. Function calling conventions are differentiated by two characteristics: the order the parameters are pushed onto the stack and whether the calling or called function is responsible for maintaining the stack.

The C calling convention specifies that arguments are placed onto the stack from right to left and that the calling function is responsible for maintaining the stack. For reference, it will be noted that the C calling convention is preferred for at least two primary reasons: native functions are not normally allowed to restore a VM stack, and native functions are often written in C.

The GCI block 208 is further responsible for ensuring that parameters specified by VM code to be passed into native code do not compromise the system. More particularly, pointers to data types and buffers are validated to ensure that they reside completely within the VM data sandbox.

More specifically, the GCI block 208 preferably includes VM wrappers for each of the called native functions. Each wrapper comprises a module that marshals parameters and return values between the “VM environment” and the “native environment.” Examples include the format of the arguments in the stack registers, alignment issues (8, 16, 32 bit aligned), etc.

As previously mentioned, the GCI block 208 is invoked using the GCALL instruction. Generally, a GCALL instruction is similar to a standard CALL function in that arguments are pushed on the stack as though a C function was going to be called. However, instead of calling a VM function with CALL, a native function is called using GCALL.

Another difference is that CALL takes the absolute address of the function to be called as a parameter, whereas GCALL takes a VM register 212 as its only parameter. The specified VM register contains a value that directs the GCI block 208 to call a particular native function. This value is referred to as a gate key and

is obtained from the symbol table 210. The gate key is also referred to herein as a “symbol” or a “magic number.”

The gate key value is an implementation specific value agreed upon by the VM engine 206 and the GCI block 208. Other than the symbol table entry, no
5 correlation between gate key values and native functions can typically be drawn. Thus, when VM code is executed in different platforms, the gate key values for each specific platform are preferably used.

The symbol table 210 provides component name, function name, argument list and gate key value (symbol) for each native function that is available to the
10 VM code. All elements in the table are preferably case-sensitive. Component names and function names adhere to the same restrictions placed on C function names.

Preferably, the standard representation of the symbol table combines the four elements into a case-sensitive ASCII string using the at-sign (@) as a
15 delimiter. A representative string format is: ComponentName@FunctionName@Arguments@GateKey. This is the gate key (magic number).

The component name is the name of an object containing member functions that operate on a particular hardware/software/firmware component of the native platform. The component name thus at least represents the logical group to which
20 a function belongs. It is desirable to group functions into components to make such more manageable, maintainable, “object-oriented” and intuitive.

The function name is simply the name of the function. Name mangling is discouraged since arguments are specified in the arguments list. The arguments list provides the number and type of arguments that the native function is
25 expecting. The GCI block 208 uses this information to validate the arguments on the VM stack for sandbox violations. Additionally, higher level languages can use the arguments list to validate that the VM code declaration of a function matches the identification of the function by the symbol table 210.

As mentioned previously, the GCI block 208 and the engine 206 have to
30 agree upon the implementation specific meaning of the gate key value. Generally, the gate key value represents a pointer to a function or an index/hash-key to an internal structure.

A preferred representation of the arguments list uses the following designations:

| | |
|----|--|
| | 1: Byte Value – no bounds checking required |
| 5 | 2: Word Value – No bounds checking required |
| | 4: Double-Word Value – no bounds checking required |
| | B: Byte Pointer – bounds checking required |
| | W: Word Pointer – bounds checking required |
| | D: Double-Word Pointer – bounds checking required |
| 10 | Z: Pointer to a NULL Terminated String – bounds checking required |
| | U: Null term to a unicode string |
| | P: Paired Argument for Buffer, Buffer Size – bounds checking required. |

15 Table 1.

It will be noted that P represents a pointer to the Buffer and is followed by a 1, 2 or 4 representing the size of the unsigned integer containing the size of the Buffer.

20 The GCI block 208 uses the arguments list to determine how many and what types of parameters need to be pushed onto the native stack from the VM stack. All accesses to the VM stack are checked to confirm that the stack lies within the VM data sandbox.

25 Additionally, if the parameter on the VM stack is a pointer, then the GCI block 208 will de-reference the pointer and confirm that the entire data element, structure or buffer pointed to lies with the VM data sandbox. The actual structure of the symbol table 210 is implementation specific.

30 Preferred VM instruction sets are represented in the tables of FIGS. 7 and 8. Generally, instructions that specify an immediate bit (noted by an I in bit 8 of the instruction coding columns) will set this bit to use the Immediate Mode of the instruction or clear this bit to use the Register Mode of the instruction. Instructions that specify granularity bits (noted by a G in bit 6 and 7 of the instruction encoding columns) determine whether the instruction will operate on the registers as a byte,

word or double-word. Registers that have an explicit bit width specified, such as the PR register, are not affected by granularity bits. Granularity encoding is preferably performed as follows:

5

| | |
|----|---------------|
| 11 | = Byte |
| 10 | = Word |
| 00 | = Double-Word |

Table 2.

10

Instructions that contain unspecified bits (noted by an X in a bit position of the instruction encoding columns) will accept any value in these bits. Such bits may be useful for encoding debug information into the code. It is recommended, however, that such bits be set to zero during normal operation.

15

Registers are encoded into instructions in three-bit groups. The three bits specify a base register. Combining a base register with the granularity bits specifies a specific register. Base register encoding is preferably performed as follows:

20

| | |
|----|------|
| 00 | = R0 |
| 01 | = R1 |
| 02 | = R2 |
| 03 | = R3 |
| 04 | = R4 |
| 05 | = R5 |
| 06 | = SP |
| 07 | = IP |

25

Table 3.

30

Immediate values are specified following the instruction word and the granularity is determined by the granularity bits unless otherwise specified in the instruction description. Immediate double-words are encoded as two words

following the instruction word with the most significant word preceding the least significant word.

Immediate words are encoded as a single word following the instruction word. Immediate bytes are encoded as a single word following the instruction
5 word. The most significant 8-bits will be ignored and should be zero filled. These bits can also be used for debug operations.

The instructions shown in FIG. 7 are divided into load/store/no operation, arithmetic and logical operational groupings. The various columns of FIG. 7 are understood as follows. The LDI instruction is a “load immediate value to register”
10 operation, which loads a designated register with an immediate value. The format of the instruction is LDI [Destination Register], [Immediate Value]. No ZCNV flags are set, and the encoding string is 0 0 1 0 0 0 0 1 G G Rd X X X. Remaining instructions shown in FIG. 6 are decoded in like manner.

The instructions shown in FIG. 8 are divided into Shift, Branch, Stack and
15 Gate Call operational groupings. By way of example, the LSL instruction is a “logical shift left” operation, which shifts the value in the destination register left the number of bits specified in the number register. Zero fills the least significant bit and discards the most significant bits. The format of the instruction is LSL [Destination Register] [32-bit Number Register]. No ZCNV flags are set, and the
20 encoding string is 1 0 0 0 0 0 0 0 G G Rd Rn. Remaining instructions shown in FIG. 8 are decoded in like manner.

FIG. 9 provides a functional block diagram representation of the PC host
140 to generally illustrate a preferred manner in which VM code is generated for the device, including the instructions to be executed and the format of the returned
25 data. Top level code is generated by the user as represented by PC code block 220. This code is preferably written in C language and identifies the particular operations desired to be carried out by the engine 206. The completed code is compiled by a VM compiler/assembler block 222, which generates corresponding object code.

30 A VM linker block 224 generates an executable (EXE) file with relocation pointers to the variously desired standard component interface functions. As shown in FIG. 10, the executable file is then provided to a VM loader block 226 which places the various gate keys in global data areas (by way of the symbol table

210). When the user subsequently launches the executable file, the file activates a VM client block 228 which manages communications with the VM server block 214 of the device 100 to execute the library of standard component interface (SCI) functions (represented by block 230).

5 It will be noted that, as discussed above, the gate keys are placed in the global data area so that, at run time, such can be loaded into the registers and called in turn. The VM engine 206 will use the symbol table 210 during gate calls to find out what parameters to load and what values will be returned. Calls involve VM functions while gate calls (GCALLs) carry out the SCI functions.

10 All of the information required by the VM engine 206, including the type of memory alignment to be carried out, are provided by the compiled PC code. The symbol table 210 basically correlates the SCI functions and the gate keys. Thus, as mentioned above the gate key operates as a pointer or an index/hash-key to an internal structure in the native code 202, thereby allowing the native function
15 (via the GCI block 208) to actually execute the function call.

 As mentioned previously, the native functions of various devices are preferably arranged as a standardized list (e.g., the Standard Component Interface or SCI functions). While not required, arranging these functions in this manner provides greater flexibility in using the same VM code across different platforms.

20 While SCI functions will tend to be device-type dependent, illustrative examples of SCI functions for a data storage device (such as 100) include sector identification, defect management, event logging, R/W functions, and servo functions, among others.

 More specifically, as those skilled in the art will appreciate, there are
25 different ways to identify specific sectors in a data storage device, including logical block address (LBA), physical block address (PBA), cylinder/head/sector (CHS), and radius/angle/length (r, θ , z). Each of these methodologies can be used to identify a particular sector address, but standardizing these across different platforms would enhance the ability to reference specific sectors in the VM code
30 regardless of the platform.

 Continuing along these lines, with regard to defect management, data storage devices that utilize SCSI based architectures typically maintain defective sectors in a sequence of lists (e.g., P, G, C and D lists). These lists indicate when

the various known defects were identified (i.e., during factory test, during field use, during a formatting operation, or other), which among other things helps to assess and predict future performance of the device. Devices that utilize other architectures (such as ATA) do not necessarily organize defective sectors in this manner, so defect list reporting may tend to be in a different format for an ATA drive than for a SCSI drive. However, using SCI function would allow a standardized way to report defects regardless of the platform, and enhance the usefulness of the VM code across different types of users.

Similar benefits can be obtained for event logging (i.e., the manner in which successively executed scripts are identified), R/W functions (error retries, reading contents of particular sectors, ECC settings, etc.), and servo functions (seek methodologies, read and write offsets, etc.).

FIGS. 11 and 12 generally illustrates the general manner in which the corresponding native function is called in response to execution of a selected virtual function. In FIG. 11, the symbol table 210 is accessed to identify the gate key (i.e., magic no. N). In FIG. 12, the GCI 208 uses the magic no. N to identify and call the corresponding native function in the native processor environment. As described above, addresses as well as descriptors are included to enable the native functions to be executed.

FIG. 13 provides a flow chart for a VIRTUAL MACHINE EMULATION routine 300, generally illustrative of steps carried out in accordance with preferred embodiments of the present invention.

At step 302, top level VM code is provided as a sequence of virtual machine instructions in a standardized format (such as previously discussed during PC code block 222). As discussed above, the instructions are processor architecture independent, and thus can be used in a number of different processors.

A symbol table (such as symbol table 210 discussed in FIGS. 5 and 9) is next generated at step 304 that maps at least selected VM symbols to corresponding native functions in the native code (block 202, FIG. 5). The symbol table, of course, is processor dependent, although the symbol table can use a standardized set of the VM instructions as a starting point (i.e., different processors will have different corresponding native functions in the table).

The VM code generated during step 302 is next modified at step 306 to substitute gate call instructions for each VM instruction that invokes a native function by the native code. This modified VM code may be therefore be processor dependent, depending upon the differences among the capabilities of the various processors, the extent to which virtual instructions can be more efficiently executed by the native processor acting as “itself,” etc. Compilation and linking steps are performed during step 306 as desired, as discussed above in FIG. 9 to prepare the code for execution by the processor.

Thereafter, as shown by step 308, the modified VM code is executed using an execution engine emulated in the memory space of the processor (such as engine 206 shown in FIGS. 5 and 9). As will be recognized from the foregoing discussion, such execution will include initiating gate calls using the gate call interface block 208 (FIGS. 5 and 9) and the symbol table 210 to invoke native functions by the processor. This allows concurrent evaluation of the execution of the processor of other functions.

FIG. 13 further serves to illustrate an advantageous manner in which the same base VM code (from step 302) is further adapted for use in a second processor having a different hardware architecture as compared to the first processor contemplated during steps 304, 306 and 308. More particularly, as shown by step 310 a second symbol table is generated to map selected VM code instructions to native instructions in the second processor. The base VM code is then modified at step 312 and executed at step 314 in similar fashion using the second processor. The process is shown to end at step 316, but it will be understood that the foregoing steps can be repeated as desired to cover any number of additional processors.

It will now be recognized that the virtual machine 204 provides significant advantages over the prior art. Developing a virtual machine processor wholly in software frees the need to tailor the machine to a particular processor hardware architecture, and eliminates the need to adapt the virtual machine as newer versions of processors or different makes and models of processors are incorporated into the product stream. This frees a device manufacturer to consider a broader range of new processors having significantly different structures as such become commercially available.

Another advantage is that development, manufacturing and field service personnel all can access the same diagnostic functions, and have data returned in the same format. This facilitates customer service and ensures that problems in the field can be more quickly resolved.

5 While preferred embodiments have been directed to data storage device (disc drive) processor-based devices, it will be recognized that this is provided for purposes of illustration and is not necessarily limiting to the scope of the claimed invention.

10 It will now be clear that preferred embodiments of the present invention are generally directed to an apparatus and method for emulating a virtual machine in the memory space of a programmable processor.

 The apparatus preferably comprises a programmable processor (such as 120) with an associated memory space (such as 122, 200) which stores native code (such as 202) comprising a sequence of native functions in a format determined by
15 a hardware architecture of the processor, the processor executing the native code to carry out at least one native function.

 The apparatus further preferably comprises a virtual machine (such as 204) within the processor memory space and concurrently emulated by the processor, the virtual machine concurrently emulated by the processor and which executes
20 virtual functions (step 314) in a standardized format determined irrespective of the hardware architecture of the processor, wherein the execution of a selected virtual function by the virtual machine causes the processor to execute a corresponding native function, and wherein the corresponding native function is executed to evaluate concurrent execution of at least one other native function.

25 Preferably, the virtual machine comprises an execution engine block (such as 206) which executes the virtual functions (such as shown in FIGS. 7 and 8), a symbol table (such as 210) that maps selected virtual functions to the corresponding native functions, and a gate call interface (GCI) block (such as 208) which calls the corresponding native functions. Preferably, execution of the native
30 function by the processor results in a return data value, and wherein the gate call interface block further operates to provide the return data value to the host device (such as 140).

It is to be understood that even though numerous characteristics and advantages of various embodiments of the present invention have been set forth in the foregoing description, together with details of the structure and function of various embodiments of the invention, this detailed description is illustrative only, and changes may be made in detail, especially in matters of structure and arrangements of parts within the principles of the present invention to the full extent indicated by the broad general meaning of the terms in which the appended claims are expressed. For example, the particular elements may vary depending on the particular application of the virtual machine without departing from the spirit and scope of the present invention.

In addition, although the embodiments described herein are directed to the emulation of a virtual machine in the memory space of a controller in a data storage device, it will be appreciated by those skilled in the art that the virtual machine and associated methodology can be used in other processor-based environments without departing from the spirit and scope of the claimed invention.